

The matching problem has no small symmetric SDP

Gábor Braun Georgia Tech	Jonah Brown-Cohen U.C. Berkeley	Arefin Huq Georgia Tech	Sebastian Pokutta* Georgia Tech
Prasad Raghavendra† U.C. Berkeley	Aurko Roy‡ Georgia Tech	Benjamin Weitz§ U.C. Berkeley	Daniel Zink Georgia Tech

December 1, 2016

Abstract

Yannakakis [1991, 1988] showed that the matching problem does not have a small symmetric linear program. Rothvoß [2014] recently proved that any, not necessarily symmetric, linear program also has exponential size. In light of this, it is natural to ask whether the matching problem can be expressed compactly in a framework such as semidefinite programming (SDP) that is more powerful than linear programming but still allows efficient optimization. We answer this question negatively for symmetric SDPs: any symmetric SDP for the matching problem has exponential size.

We also show that an $O(k)$ -round Lasserre SDP relaxation for the asymmetric metric traveling salesperson problem yields at least as good an approximation as any *symmetric* SDP relaxation of size n^k .

The key technical ingredient underlying both these results is an upper bound on the degree needed to derive polynomial identities that hold over the space of matchings or traveling salesperson tours.

*Research reported in this paper was partially supported by NSF CAREER award CMMI-1452463 and NSF grant CCF-1255900.

†This work is supported by the National Science Foundation under the NSF CAREER Award CCF-1407779, and the Sloan Fellowship from the Alfred P. Sloan Foundation.

‡Research reported in this paper was partially supported by NSF grant CMMI-1333789.

§This work was supported by the National Science Foundation Graduate Research Fellowship Program award DGE 1106400

1 Introduction

In his seminal work, Yannakakis [1991, 1988] showed that any symmetric linear program for the matching problem has exponential size. Rothvoß [2014] recently showed that one can drop the symmetry requirement: any linear program for the matching problem has exponential size. Since it is possible to optimize over matchings in polynomial time, it follows that there is a gap between problems that have small linear formulations and problems that allow efficient optimization.

In light of this gap, it is reasonable to ask whether semidefinite programming (SDP) can characterize all problems that allow efficient optimization. Semidefinite programs generalize linear programs and can be solved efficiently both in theory and practice (see Vandenberghe and Boyd [1996]). SDPs are the basis of some of the best algorithms currently known, for example the approximation of Goemans and Williamson [1995] for MAX CUT .

Following prior work (see for example Gouveia et al. [2011]) we define the size of an SDP formulation as the dimension of the psd cone from which the polytope can be obtained as an affine slice. Some recent work has shown limits to the power of small SDPs. Briët et al. [2013, 2015] nonconstructively give an exponential lower bound on the size of SDP formulations for most 0/1 polytopes. Lee et al. [2015] give an exponential lower bound for solving the traveling salesperson problem (TSP) and approximating MAX 3-SAT. However the question of whether the matching problem has a small SDP remains open. We give a partial negative answer to this question by proving the analog of Yannakakis's result for semidefinite programs:

Theorem 1.1. *Any symmetric SDP for the matching problem has exponential size.*

As we explain below, the main challenge we faced in obtaining this result was to develop machinery to handle the nontrivial structure of the solution space of matchings.

Using a similar argument, we also show that for the asymmetric metric traveling salesperson problem the optimal symmetric semidefinite formulation of a given size is essentially achieved by the respective level of the Lasserre hierarchy.

Related work

Bounding the size of general linear programming formulations for a given problem was initiated by the seminal paper of Yannakakis [1991, 1988]. In Yannakakis's model, a general linear program for say the perfect matching polytope $\text{PM}(n)$ consists of a higher-dimensional polytope $Q \in \mathbb{R}^D$ and a projection π such that $\pi(Q) = \text{PM}(n)$. The size of the linear program is measured as the number of inequalities required to define the polytope Q .

Yannakakis [1991] characterized the size of linear programming formulations in terms of the non-negative rank of an associated matrix known as the *slack matrix*. Using this characterization, Yannakakis showed that any *symmetric* linear program for the matching problem or traveling salesperson problem requires exponential size. Roughly speaking, a linear program for the matching problem is symmetric if for every permutation σ of the vertices in the corresponding graph, there is a permutation $\tilde{\sigma}$ of the coordinates in \mathbb{R}^D that leaves the linear program (and thus the polytope Q) unchanged.

A natural question that came out of the work of Yannakakis is whether dropping the symmetry requirement helps much. Kaibel et al. [2010] showed that dropping the symmetry requirement can mean the difference between polynomial and superpolynomial size linear extended formulations

for matchings with $\lfloor \log n \rfloor$ edges in K_n , while Goemans [2015] and Pashkovich [2014] showed that it can mean the difference between subquadratic and quadratic size linear extended formulations for the permutohedron. Nonetheless, Fiorini et al. [2012, 2015] and Rothvoß [2014] answered this question negatively for the TSP and matching problems respectively: any linear extended formulation of either problem, symmetric or not, has exponential size. In particular, Fiorini et al. [2012, 2015] established a $2^{\Omega(\sqrt{n})}$ lower bound on any linear formulation for the TSP, while Rothvoß [2014] later established a $2^{\Omega(n)}$ lower bound on any linear formulation for matching, which by a known reduction implies a $2^{\Omega(n)}$ lower bound for the TSP. From a computational standpoint, these are strong lower bounds against solving the TSP or matching problem exactly via small linear programs. Subsequently, the framework of Yannakakis has been generalized towards showing lower bounds even for approximating combinatorial optimization problems in Braun et al. [2012, 2015a], Chan et al. [2013], Braverman and Moitra [2013], Bazzi et al. [2015].

For the class of maximum constraint satisfaction problems (MaxCSPs), Chan et al. [2013] established a connection between lower bounds for general linear programs and lower bounds against an explicit linear program, namely the hierarchy of Sherali and Adams [1990]. Using that connection, Chan et al. [2013] showed that for every constant d and for every MaxCSP, the d -round Sherali-Adams LP relaxation yields at least as good an approximation as any LP relaxation of size $n^{d/2}$. By appealing to lower bounds on Sherali-Adams relaxations of MaxCSPs in literature, they give super-polynomial lower bounds for MAX 3-SAT and other MaxCSPs.

Given the general LP lower bounds, it is natural to ask whether the situation is different for SDP relaxations. Building on the approach of Chan et al. [2013], Lee et al. [2014] showed that for the class of MaxCSPs, for every constant d , the d -round Lasserre SDP relaxation yields at least as good an approximation as any symmetric SDP relaxation of size $n^{d/2}$. In light of known lower bounds for Lasserre SDP relaxations of MAX 3-SAT, this yields a corresponding lower bound for approximating MAX 3-SAT. In a recent advance, Lee et al. [2015] show an exponential lower bound even for *asymmetric* SDP relaxations of the TSP.

Contribution

We show that there is no small symmetric SDP for the matching problem. Our result is an SDP analog of the result in Yannakakis [1991, 1988] that ruled out a small symmetric LP for the matching problem. Specifically, we show:

Theorem 1.2. *There exists an absolute constant $\alpha > 0$ such that for every $\varepsilon \in [0, 1)$, every symmetric SDP relaxation approximating the perfect matching problem within a factor $1 - \frac{\varepsilon}{n-1}$ has size at least $2^{\alpha n}$.*

To prove this we show that if the matching problem has a small symmetric SDP relaxation, then there is a low-degree sum of squares refutation of the existence of a perfect matching in an odd clique, which contradicts a result by Grigoriev [2001].

The key technical obstacle in adapting the MaxCSP argument to the matching problem is the non-trivial algebraic structure of the underlying solution space (the space of all perfect matchings). A multilinear polynomial is zero over the solution space of a MaxCSP (the Boolean hypercube $\{0, 1\}^n$) if and only if all the coefficients of the polynomial are zero, which is trivial to test. In contrast, simply testing whether a multilinear polynomial is zero over all perfect matchings is non-trivial. Nonetheless, in a key lemma we show that every multilinear polynomial F that is identically zero

over perfect matchings can be certified as such via a derivation of degree $2 \cdot \deg(F) - 1$, starting from the linear and quadratic constraints that define the space of perfect matchings.

Our second result concerns the asymmetric metric traveling salesperson problem, which is a restriction of the TSP to the case where edge costs obey the triangle equality but are not necessarily symmetric, so that the cost from u to v may not equal the cost from v to u . We show that Lasserre SDP relaxations are more or less optimal among all symmetric SDP relaxations for approximating the asymmetric metric traveling salesperson problem. The precise statement follows.

Theorem 1.3. *For every constant $\rho > 0$, if there exists a symmetric SDP relaxation of size $r < \sqrt{\binom{2n}{k}} - 1$ which achieves a ρ -approximation for asymmetric metric TSP instances on $2n$ vertices, then the $(2k-1)$ -round Lasserre relaxation achieves a ρ -approximation for asymmetric metric TSP instances on n vertices.*

2 Symmetric SDP formulations

In this section we define a framework for symmetric semidefinite programming formulations and show that a symmetric SDP formulation implies a symmetric sum of squares representation over a small basis. Our framework extends the one in [Braun et al. \[2015b\]](#) with a symmetry condition; see also [Lee et al. \[2014\]](#).

We first introduce some notation we will use. Let $[n]$ denote the set $\{1, \dots, n\}$. Let \mathbb{S}_+^r denote the cone of $r \times r$ real symmetric positive semidefinite (psd) matrices. Let $\mathbb{R}[x]$ denote the set of polynomials in n real variables $x = (x_1, \dots, x_n)$ with real coefficients. For a set $\mathcal{H} \subseteq \mathbb{R}[x]$ let $\langle \mathcal{H} \rangle$ denote the vector space spanned by \mathcal{H} and let $\langle \mathcal{H} \rangle_I$ denote the ideal generated by \mathcal{H} . (Recall that a polynomial *ideal* in a polynomial ring R is a set that is closed under addition of polynomials in the ideal and closed under multiplication by polynomials in the ring.)

Suppose group G acts on a set X . The (left) action of $g \in G$ on $x \in X$ is denoted $g \cdot x$. Recall that the *orbit* of $x \in X$ is $\{g \cdot x \mid g \in G\}$ while the *stabilizer* of x is $\{g \in G \mid g \cdot x = x\}$. Let A_n denote the alternating group on n letters (the set of even permutations of $[n]$).

We now present our SDP formulation framework. We restrict ourselves to maximization problems even though the framework extends to minimization problems. A *maximization problem* $\mathcal{P} = (\mathcal{S}, \mathcal{F})$ consists of a set \mathcal{S} of feasible solutions and a set \mathcal{F} of objective functions. Suppose we are given two functions $\tilde{C}: \mathcal{F} \rightarrow \mathbb{R}$ and $\tilde{S}: \mathcal{F} \rightarrow \mathbb{R}$. We say an algorithm (\tilde{C}, \tilde{S}) -approximately solves \mathcal{P} if for all $f \in \mathcal{F}$ with $\max_{s \in \mathcal{S}} f(s) \leq \tilde{S}(f)$ it computes $\tilde{f} \in \mathbb{R}$ satisfying $\max_{s \in \mathcal{S}} f(s) \leq \tilde{f} \leq \tilde{C}(f)$. We will refer to \tilde{C} and \tilde{S} as the *approximation guarantees*.

Example 2.1. Suppose we are given a polytope P with non-redundant inner and outer descriptions:

$$P = \text{conv}(V) = \left\{ x \mid a_j x \leq b_j, j \in [m] \right\}.$$

Let us define $f_j(x) := b_j - a_j x$ for each $j \in [m]$. We can now associate a maximization problem with this polytope by setting $\mathcal{S} = V$ and $\mathcal{F} = \{f_j \mid j \in [m]\}$, so that each vertex is a solution and the slack with respect to each facet is a function. In order to recover the polytope exactly we would set

$$\tilde{C}(f) = \tilde{S}(f) = \max_{x \in P} f(x) = 0$$

for all $f \in \mathcal{F}$.

Let G be a group with associated actions on \mathcal{S} and \mathcal{F} . The problem \mathcal{P} is G -symmetric if the group action satisfies the compatibility constraint $(g \cdot f)(g \cdot s) = f(s)$. For a G -symmetric problem we require G -symmetric approximation guarantees: $\tilde{C}(g \cdot f) = \tilde{C}(f)$ and $\tilde{S}(g \cdot f) = \tilde{S}(f)$ for all $f \in \mathcal{F}$ and $g \in G$.

We now define the notion of a semidefinite programming formulation of a maximization problem.

Definition 2.2 (SDP formulation for \mathcal{P}). Let $\mathcal{P} = (\mathcal{S}, \mathcal{F})$ be a maximization problem with approximation guarantees \tilde{C}, \tilde{S} . A (\tilde{C}, \tilde{S}) -approximate SDP formulation of \mathcal{P} of size d consists of a linear map $\mathcal{A}: \mathbb{S}_+^d \rightarrow \mathbb{R}^k$ and $b \in \mathbb{R}^k$ together with

1. *Feasible solutions:* an $X^s \in \mathbb{S}_+^d$ with $\mathcal{A}(X^s) = b$ for all $s \in \mathcal{S}$, i.e., the SDP $\{X \in \mathbb{S}_+^d \mid \mathcal{A}(X) = b\}$ is a relaxation of $\text{conv} \{X^s \mid s \in \mathcal{S}\}$,
2. *Objective functions:* an affine function $w^f: \mathbb{S}_+^d \rightarrow \mathbb{R}$ satisfying $w^f(X^s) = f(s)$ for all $f \in \mathcal{F}$ with $\max_{s \in \mathcal{S}} f(s) \leq \tilde{S}(f)$ and all $s \in \mathcal{S}$, i.e., the linearizations are exact on solutions, and
3. *Achieving guarantee:* $\max \{w^f(X) \mid \mathcal{A}(X) = b, X \in \mathbb{S}_+^d\} \leq \tilde{C}(f)$ for all $f \in \mathcal{F}$ with $\max_{s \in \mathcal{S}} f(s) \leq \tilde{S}(f)$.

If G is a group, \mathcal{P} is G -symmetric, and G acts on \mathbb{S}_+^d , then an SDP formulation of \mathcal{P} with symmetric approximation guarantees \tilde{C}, \tilde{S} is G -symmetric if it additionally satisfies the compatibility conditions for all $g \in G$:

1. *Action on solutions:* $X^{g \cdot s} = g \cdot X^s$ for all $s \in \mathcal{S}$.
2. *Action on functions:* $w^{g \cdot f}(g \cdot X) = w^f(X)$ for all $f \in \mathcal{F}$ with $\max_{s \in \mathcal{S}} f(s) \leq \tilde{S}(f)$.
3. *Invariant affine space:* $\mathcal{A}(g \cdot X) = \mathcal{A}(X)$.

A G -symmetric SDP formulation is G -coordinate-symmetric if the action of G on \mathbb{S}_+^d is by permutation of coordinates: that is, there is an action of G on $[d]$ with $(g \cdot X)_{ij} = X_{g^{-1}i, g^{-1}j}$ for all $X \in \mathbb{S}_+^d$, $i, j \in [d]$ and $g \in G$.

We now turn a G -coordinate-symmetric SDP formulation into a symmetric sum of squares representation over a small set of basis functions.

Lemma 2.3 (Sum of squares for a symmetric SDP formulation). *If a G -symmetric maximization problem $\mathcal{P} = (\mathcal{S}, \mathcal{F})$ admits a G -coordinate-symmetric (\tilde{C}, \tilde{S}) -approximate SDP formulation of size d , then there is a set \mathcal{H} of at most $\binom{d+1}{2}$ functions $h: \mathcal{S} \rightarrow \mathbb{R}$ such that for any $f \in \mathcal{F}$ with $\max f \leq \tilde{S}(f)$ we have $\tilde{C}(f) - f = \sum_j h_j^2 + \mu_f$ for some $h_j \in \langle \mathcal{H} \rangle$ and constant $\mu_f \geq 0$. Furthermore the set \mathcal{H} is invariant under the action of G given by $(g \cdot h)(s) = h(g^{-1} \cdot s)$ for $g \in G$, $h \in \mathcal{H}$ and $s \in \mathcal{S}$.*

Proof. For any psd matrix M let \sqrt{M} denote the unique psd matrix with $\sqrt{M}^2 = M$. Note that $\sqrt{M} \sqrt{M}^\top = M$ also, since \sqrt{M} is symmetric.

Let $\mathcal{A}, b, \{X^s \mid s \in \mathcal{S}\}, \{w^f \mid f \in \mathcal{F}\}$ comprise a G -coordinate-symmetric SDP formulation of size d . We define the set $\mathcal{H} := \{h_{ij} \mid i, j \in [d]\}$ via $h_{ij}(s) := \sqrt{X^s}_{ij}$. By the action of G and the uniqueness of the square root, we have $g \cdot h_{ij} = h_{g \cdot i, g \cdot j}$, so \mathcal{H} is G -symmetric. As $h_{ij} = h_{ji}$, the set \mathcal{H} has at most $\binom{d+1}{2}$ elements.

By standard strong duality arguments as in [Braun et al. \[2015b\]](#), for every $f \in \mathcal{F}$ with $\max f \leq \tilde{S}(f)$, there is a $U^f \in \mathbb{S}_+^d$ and $\mu_f \geq 0$ such that for all $s \in \mathcal{S}$,

$$\tilde{C}(f) - f(s) = \text{Tr}[U^f X^s] + \mu_f.$$

Again by standard arguments the trace can be rewritten as a sum of squares:

$$\text{Tr}[U^f X^s] = \text{Tr}\left[\left(\sqrt{U^f} \sqrt{X^s}\right)^\top \left(\sqrt{U^f} \sqrt{X^s}\right)\right] = \sum_{i,j \in [d]} \left(\sum_{k \in [d]} \sqrt{U^f}_{ik} \cdot \sqrt{X^s}_{kj} \right)^2.$$

Therefore $\tilde{C}(f) - f = \sum_{i,j \in [d]} \left(\sum_{k \in [d]} \sqrt{U^f}_{ik} \cdot h_{kj} \right)^2 + \mu_f$, as claimed. \square

3 The perfect matching problem

We now present the *perfect matching problem* $\text{PM}(n)$ as a maximization problem in the framework of [Section 2](#) and show that any symmetric SDP formulation has exponential size.

Let n be an even positive integer, and let K_n denote the complete graph on n vertices. The feasible solutions of $\text{PM}(n)$ are all the perfect matchings M on K_n . The objective functions f_F are indexed by the edge sets F of K_n : for each $F \subseteq \binom{[n]}{2}$ we define $f_F(M) := |M \cap F|$. For approximation guarantees we use $\tilde{S}(f) := \max f$ and $\tilde{C}(f) := \max f + \varepsilon/2$ for some fixed $0 \leq \varepsilon < 1$ as in [Braun and Pokutta \[2015a\]](#); see also [Braun and Pokutta \[2015b\]](#) for a more in-depth discussion.

3.1 Symmetric functions on matchings are juntas

In this section we show that functions on perfect matchings with high symmetry are actually *juntas*: they depend only on the edges of a small vertex set. The key is the following lemma stating that perfect matchings coinciding on a vertex set belong to the same orbit of the pointwise stabilizer of the vertex set. Let A_n denote the alternating group on n letters, and for any subset $X \subseteq [n]$ let $A(X)$ denote the alternating group that operates on the elements of X and fixes the remaining elements of $[n]$. For any set $W \subseteq [n]$ let $E[W]$ denote the edges of K_n with both endpoints in W .

Lemma 3.1. *Let $S \subseteq [n]$ with $|S| < n/2$ and let M_1 and M_2 be perfect matchings in K_n . If $M_1 \cap E[S] = M_2 \cap E[S]$ then there exists $\sigma \in A([n] \setminus S)$ such that $\sigma \cdot M_1 = M_2$.*

Proof. Let $\delta(S)$ denote the edges with exactly one endpoint in S . There are three kinds of edges: those in $E[S]$, those in $\delta(S)$, and those disjoint from S . We construct σ to handle each type of edge, then fix σ to be even.

To handle the edges in $E[S]$ we set σ to the identity on S , since $M_1 \cap E[S] = M_2 \cap E[S]$.

To handle the edges in $\delta(S)$ we note that for each edge $(s, v) \in M_1$ with $s \in S$ and $v \notin S$ there is a unique edge $(s, w) \in M_2$ with $w \notin S$. We extend σ to map v to w for each such s .

To handle the edges disjoint from S , we again use the fact that M_1 and M_2 are perfect matchings, so the number of edges in each that are disjoint from S is the same. We extend σ to be an arbitrary bijection on those edges.

We now show that we can choose σ to be even. Since $|S| < n/2$ there is an edge $(u, v) \in M_2$ disjoint from S . Let $\tau_{u,v}$ denote the transposition of u and v and let $\sigma' := \tau_{u,v} \circ \sigma$. We have $\sigma' \cdot M_1 = \sigma \cdot M_1 = M_2$, and either σ or σ' is even. \square

We also need the following lemma, which has been used extensively for symmetric linear extended formulations. See references Yannakakis [1988, 1991], Kaibel et al. [2010], Braun and Pokutta [2012], Lee et al. [2014] for examples.

Lemma 3.2 ([Dixon and Mortimer, 1996, Theorems 5.2A and 5.2B]). *Let $n \geq 10$ and let $G \leq A_n$ be a group. If $|A_n : G| < \binom{n}{k}$ for some $k < n/2$, then there is a subset $W \subseteq [n]$ such that $|W| < k$, W is G -invariant, and $A([n] \setminus W)$ is a subgroup of G .*

We now formally state and prove the claim about juntas:

Proposition 3.3. *Let $n \geq 10$, let $k < n/2$ and let \mathcal{H} be an A_n -symmetric set of functions on the set of perfect matchings of K_n of size less than $\binom{n}{k}$. Then for every $h \in \mathcal{H}$ there is a vertex set $W \subseteq [n]$ of size less than k such that h depends only on the (at most $\binom{k-1}{2}$) edges in W .*

Proof. Let $h \in \mathcal{H}$, let $\text{Stab}(h)$ denote the stabilizer of h , and let $\text{Orb}(h)$ denote the orbit of h . Since \mathcal{H} is A_n -symmetric we have $|\text{Orb}(h)| < \binom{n}{k}$. By the orbit-stabilizer theorem it follows that $|A_n : \text{Stab}(h)| < \binom{n}{k}$. Applying Lemma 3.2 to the stabilizer of h , we obtain a subset $W \subseteq [n]$ of size less than k such that h is stabilized by $A([n] \setminus W)$, i.e.,

$$h(M) = (g \cdot h)(M) = h(g^{-1} \cdot M)$$

for all $g \in A([n] \setminus W)$.

Therefore for every perfect matching M the function h is constant on the $A([n] \setminus W)$ -orbit of M . As the orbit is determined by $M \cap E[W]$ by Lemma 3.1, so is the function value $h(M)$. Therefore h depends only on the edges in $E[W]$. \square

3.2 The matching polynomials

A key step in proving our lower bound is obtaining low-degree derivations of approximation guarantees for objective functions of $\text{PM}(n)$. Therefore we start with a standard representation of functions as polynomials. We define the *matching constraint polynomials* \mathcal{P}_n as:

$$\begin{aligned} \mathcal{P}_n := & \{x_{uv}x_{uw} \mid u, v, w \in [n] \text{ distinct}\} \\ & \cup \left\{ \sum_{u \in [n], u \neq v} x_{uv} - 1 \mid v \in [n] \right\} \\ & \cup \{x_{uv}^2 - x_{uv} \mid u, v \in [n] \text{ distinct}\}. \end{aligned} \tag{3.1}$$

We observe that the ring of real valued functions on perfect matchings is isomorphic to $\mathbb{R}[\{x_{uv}\}_{\{u,v\} \in \binom{[n]}{2}}]/\langle \mathcal{P}_n \rangle_I$ with x_{uv} representing the indicator function of the edge uv being contained in a perfect matching. Intuitively, under this representation the vanishing of the first set of polynomials ensures that no vertex is matched more than once, the vanishing of the second set ensures that each vertex is matched, and the vanishing of the third set ensures that each edge coordinate is 0-1 valued.

Now we formulate low-degree derivations. Let \mathcal{P} denote a set of polynomials in $\mathbb{R}[x]$. For polynomials F and G , we write $F \simeq_{(\mathcal{P}, d)} G$, or F is congruent to G from \mathcal{P} in degree d , if and only if there exist polynomials $\{q(p) : p \in \mathcal{P}\}$ such that

$$F + \sum_{p \in \mathcal{P}} q(p) \cdot p = G$$

and $\max_p \deg(q(p) \cdot p) \leq d$. We often drop the dependence on \mathcal{P} when it is clear from context. We shall write $F \equiv G$ for two polynomials F and G defining the same function on perfect matchings, i.e., $F - G \in \langle \mathcal{P}_n \rangle_I$. (Note that as \mathcal{P}_n contains $x_{uv}^2 - x_{uv}$ for all variables x_{uv} , the ideal generated by \mathcal{P}_n is automatically radical.)

3.3 Deriving that symmetrized polynomials are constant

Averaging any polynomial on matchings over the symmetric group gives a constant. In this section we show that this fact has a low-degree derivation.

For a partial matching M , let $x_M := \prod_{e \in M} x_e$ denote the product of edge variables for the edges in M . The first step is to reduce every polynomial to a linear combination of the x_M .

Lemma 3.4. *For every polynomial F there is a polynomial F' with $\deg F' \leq \deg F$ and $F \simeq_{(\mathcal{P}_n, \deg F)} F'$, where all monomials of F' have the form x_M for some partial matching M .*

Proof. It suffices to prove the lemma when F is a monomial. Let $F = \prod_{e \in A} x_e^{k_e}$ for a set A of edges with multiplicities $k_e \geq 1$. From $x_e^2 \simeq_2 x_e$ it follows that $x_e^k \simeq_k x_e$ for all $k \geq 1$, hence $F \simeq_{\deg F} \prod_{e \in A} x_e$. If A is a partial matching we are done, otherwise there are distinct $e, f \in A$ with a common vertex, hence $x_e x_f \simeq_2 0$ and $F \simeq_{\deg F} 0$. \square

Lemma 3.5. *For any partial matching M on $2d$ vertices and a vertex a not covered by M , we have*

$$x_M \simeq_{(\mathcal{P}_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, u\} \\ u \in K_n \setminus (M \cup \{a\})}} x_{M_1}. \quad (3.2)$$

Proof. We use the generators $\sum_u x_{au} - 1$ to add variables corresponding to edges at a , and then use $x_{au} x_{uv}$ to remove monomials not corresponding to a partial matching:

$$x_M \simeq_{(\mathcal{P}_n, d+1)} x_M \sum_{u \in K_n} x_{au} \simeq_{(\mathcal{P}_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, u\} \\ u \in K_n \setminus (M \cup \{a\})}} x_{M_1}.$$

\square

This leads to a similar congruence using all containing matchings of a larger size:

Lemma 3.6. *For any partial matching M of $2d$ vertices and $d \leq k \leq n/2$, we have*

$$x_M \simeq_{(\mathcal{P}_n, k)} \frac{1}{\binom{n/2-d}{k-d}} \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'} \quad (3.3)$$

Proof. We use induction on $k - d$. The start of the induction is with $k = d$, when the sides of (3.3) are actually equal. If $k > d$, let a be a fixed vertex not covered by M . Applying Lemma 3.5 to M and a followed by the inductive hypothesis gives:

$$x_M \simeq_{(\mathcal{P}_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, u\} \\ u \in K_n \setminus (M \cup \{a\})}} x_{M_1} \simeq_{(\mathcal{P}_n, k)} \frac{1}{\binom{n/2-d-1}{k-d-1}} \sum_{\substack{M' \supset M_1 \\ |M'|=k \\ M_1 = M \cup \{a, u\} \\ u \in K_n \setminus (M \cup \{a\})}} x_{M'}.$$

Averaging over all vertices a not covered by M , we obtain:

$$x_M \simeq_{(\mathcal{P}_n, k)} \frac{1}{n-2d} \frac{1}{\binom{n/2-d-1}{k-d-1}} \sum_{\substack{M' \supset M_1 \\ |M'|=k \\ M_1 = M \cup \{a, u\} \\ a, u \in K_n \setminus M}} x_{M'} = \frac{1}{n-2d} \frac{1}{\binom{n/2-d-1}{k-d-1}} 2(k-d) \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'} = \frac{1}{\binom{n/2-d}{k-d}} \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'}.$$

where in the second step the factor $2(k-d)$ accounts for the number of ways to choose a and u . \square

We are now ready to state and prove the claim about symmetrized polynomials:

Lemma 3.7. *For any polynomial F , there is a constant c_F with $\sum_{\sigma \in S_n} \sigma F \simeq_{(\mathcal{P}_n, \deg F)} c_F$.*

Proof. Given Lemma 3.4, it suffices to prove the claim for $F = x_M$ for some partial matching M . Note that if $|M| = k$ then (using the orbit-stabilizer theorem) the size of the stabilizer of M is $2^k k!(n-2k)!$. Now apply Lemma 3.6 with $d = 0$:

$$\sum_{\sigma \in S_n} \sigma x_M = 2^k k!(n-2k)! \sum_{M': |M'|=k} x_{M'} \simeq_k 2^k k!(n-2k)! \binom{n/2}{k}.$$

\square

3.4 Low-degree certificates for matching ideal membership

In this section we present a crucial part of our argument, namely that every degree- d polynomial that is identically zero over perfect matchings has a derivation of this fact whose degree is $O(d)$.

The following lemma will allow us to apply induction:

Lemma 3.8. *If L is a polynomial with $L \simeq_{(\mathcal{P}_{n-2}, d)} 0$ for some d , and a, b are the two additional vertices in K_n , then $Lx_{ab} \simeq_{(\mathcal{P}_n, d+1)} 0$.*

Proof. It is enough to prove the claim for $L \in \mathcal{P}_{n-2}$. For $L = x_e^2 - x_e$ and $L = x_{uv}x_{uw}$ the claim is trivial since $L \in \mathcal{P}_n$ also. The remaining case is $L = \sum_{u \in K_{n-2}} x_{uv} - 1$ for some $v \in K_{n-2}$. Then

$$Lx_{ab} = \left(\sum_{u \in K_{n-2}} x_{uv} - 1 \right) x_{ab} = \left(\sum_{u \in K_n} x_{uv} - 1 \right) x_{ab} - x_{av}x_{ab} - x_{bv}x_{ab} \simeq_{d+1} 0.$$

The degree of the derivation is at most $d+1$ since we can simply multiply the degree- d derivation for $L \simeq 0$ by x_{ab} . \square

We now show that any $F \in \langle \mathcal{P}_n \rangle_I$ can be generated by low-degree coefficients from \mathcal{P}_n :

Theorem 3.9. *For every polynomial $F \in \mathbb{R}[x_{uv}]_{\{u, v\} \in \binom{[n]}{2}}$, if $F \in \langle \mathcal{P}_n \rangle_I$ then $F \simeq_{(\mathcal{P}_n, 2 \deg F - 1)} 0$.*

Proof. We use induction on the degree d of F . If $d = 0$ then $F = 0$ and the statement holds trivially. (Note that \simeq_{-1} is just equality.) The case $d = 1$ rephrased means that the affine space spanned by the characteristic vectors of all perfect matchings is defined by the $\sum_v x_{uv} - 1$ for all vertices u . This follows from Edmonds's description of the perfect matching polytope by linear inequalities in Edmonds [1965].

For the case $d \geq 2$ we first prove the following claim:

Claim. If $F \in \langle \mathcal{P}_n \rangle_I$ is a degree- d polynomial and $\sigma \in S_n$ is a permutation of vertices, then

$$F \simeq_{(\mathcal{P}_n, 2d-1)} \sigma F.$$

We use induction on the degree. If $d = 0$ or $d = 1$ the claim follows from the corresponding cases $d = 0$ and $d = 1$ of the theorem. For $d \geq 2$ it is enough to prove the claim when σ is a transposition of two vertices a and u . Note that in $F - \sigma F$ all monomials which are independent of both a and u cancel:

$$F - \sigma F = \sum_{e: a \in e \text{ or } u \in e} L_e x_e \quad (3.4)$$

where each L_e has degree at most $d - 1$. We now show that every summand is congruent to a sum of monomials containing edges incident to both a and u . For example, for $e = \{a, b\}$ in (3.4) we apply the generator $\sum_v x_{uv} - 1$ to find:

$$L_{ab} x_{ab} \simeq_{d+1} L_{ab} x_{ab} \sum_v x_{uv} \simeq_{d+1} \sum_v L_{ab} x_{ab} x_{uv}.$$

Therefore

$$F - \sigma F \simeq_{d+1} \sum_{bv} L'_{bv} x_{ab} x_{uv}$$

for some polynomials L'_{bv} of degree at most $d - 1$. We may assume that L'_{bv} does not contain variables x_e with e incident to a, b, u, v , as these can be removed using generators like $x_{ab} x_{ac}$ or $x_{ab}^2 - x_{ab}$. Moreover, it can be checked that L'_{bv} is zero on all perfect matchings containing $\{a, b\}$ and $\{u, v\}$. By induction, $L'_{bv} \simeq_{(\mathcal{P}_{n-4}, 2d-3)} 0$ (identifying K_{n-4} with the graph $K_n \setminus \{a, b, u, v\}$), from which $L'_{bv} \simeq_{(\mathcal{P}_n, 2d-1)} 0$ follows by two applications of Lemma 3.8. (The special case $a = v, b = u$ is also handled by induction and one application of Lemma 3.8.) This concludes the proof of the claim.

We now apply the claim followed by Lemma 3.7:

$$F \simeq_{2d-1} \frac{1}{n!} \sum_{\sigma \in S_n} \sigma F \simeq_d \frac{c_F}{n!}$$

for a constant c_F . As $F \in \langle \mathcal{P}_n \rangle_I$, it must be that $c_F = 0$, and therefore $F \simeq_{2d-1} 0$. \square

3.5 The main theorem

We now have all the ingredients to prove our main theorem. Note that the alternating group A_n acts naturally on $\text{PM}(n)$ via permutation of vertices. Recall that we set $\tilde{S}(f) := \max f$ and $\tilde{C}(f) := \max f + \varepsilon/2$, where the functions f are indexed by edge set and ε is a parameter. It follows that the guarantees \tilde{C}, \tilde{S} are A_n -symmetric in the sense defined in Section 2. Our main theorem is an exponential lower bound on the size of any A_n -coordinate-symmetric SDP extension of $\text{PM}(n)$.

Theorem 3.10 (Main). *There exists an absolute constant $\alpha > 0$ such that for all even n and every $0 \leq \varepsilon < 1$, every A_n -coordinate-symmetric SDP extended formulation approximating the perfect matching problem $\text{PM}(n)$ within a factor of $1 - \varepsilon/(n-1)$ has size at least $2^{\alpha n}$.*

Proof. Fix an even integer $n \geq 10$ and let $k = \lceil \beta n \rceil$ for some small enough constant $0 < \beta < 1/2$ chosen later. Suppose for a contradiction that $\text{PM}(n)$ admits a symmetric SDP extended formulation of size $d < \sqrt{\binom{n}{k}} - 1$.

Let m equal $n/2$ or $n/2 - 1$, whichever is odd. Let $S = [m]$ and let $T = \{m + 1, \dots, 2m\}$. If $m = n/2$ then let $U = \{2m + 1, 2m + 2\}$, otherwise let $U = \emptyset$. Note that $S \cup T \cup U = [n]$ and $|S| = |T| = m = \Theta(n)$. Consider the objective function for the set of edges $E[S]$, namely $f_{E[S]}(M) := |M \cap E[S]|$. Since $|S|$ is odd we have $\max f_{E[S]} = (|S| - 1)/2$, from which we obtain:

$$f(x) \stackrel{\text{def}}{=} \tilde{C}(f_{E[S]}) - f_{E[S]}(x) = \frac{|S| - 1}{2} + \frac{\varepsilon}{2} - \sum_{u,v \in S} x_{uv} \equiv \frac{1}{2} \sum_{u \in S, v \in T \cup U} x_{uv} - \frac{1 - \varepsilon}{2}. \quad (3.5)$$

By Lemma 2.3, as $\binom{d+1}{2} < \binom{n}{k}$, there is a constant $\mu_f \geq 0$ and an A_n -symmetric set \mathcal{H} of functions of size at most $\binom{n}{k}$ on the set of perfect matchings with

$$f \equiv \sum_g g^2 + \mu_f \quad \text{with each } g \in \langle \mathcal{H} \rangle.$$

By Proposition 3.3, every $h \in \mathcal{H}$ depends only on the edges within a vertex set of size less than k , and hence can be represented by a polynomial of degree less than $k/2$ over perfect matchings. As the g are linear combinations of the $h \in \mathcal{H}$, they can also be represented by polynomials of degree less than $k/2$, which we assume for the rest of the proof.

Applying Theorem 3.9 with (3.5), we conclude

$$\frac{1}{2} \sum_{u \in S, v \in T \cup U} x_{uv} - \frac{1 - \varepsilon}{2} \simeq_{(\mathcal{P}_n, 2k-1)} \sum_g g^2 + \mu_f.$$

We now apply the following substitution: set $x_{2m+1, 2m+2} := 1$ if U is not empty, set $x_{u+m, v+m} := x_{uv}$ for each $uv \in E[S]$, and set $x_{uv} := 0$ otherwise. Intuitively, the substitution ensures that U is matched, ensures the matching on T is identical to the matching on S , and ensures every edge is entirely within S , T , or U . The main point is that the substitution maps every polynomial in \mathcal{P}_n either to 0 or into \mathcal{P}_m .

Applying this substitution we obtain a new polynomial identity on the variables $\{x_{uv}\}_{\{u,v\} \in \binom{[n]}{2}}$:

$$-\frac{1 - \varepsilon}{2} \simeq_{(\mathcal{P}_m, 2k-1)} \sum_g g^2 + \mu_f. \quad (3.6)$$

This equation is a sum of squares proof that an odd clique of size m cannot have a perfect matching. To complete our argument we appeal to a theorem from Grigoriev [2001] which shows that any such proof must have high degree. Since the degree of the proof in (3.6) is $2k - 1$, our conclusion will be that k must be large.

The theorem from Grigoriev [2001] uses different terminology from what we have developed here. It is phrased in terms of Positivstellensatz Calculus ($PC_>$) proof systems and the MOD₂ principle. We first present the theorem as originally stated and then relate it to our setting.

Theorem 3.11 ([Grigoriev, 2001, Corollary 2]). *The degree of any $PC_>$ refutation of MOD_2^k is greater than $\Omega(k)$.*

The MOD_p^k principle states that it is not possible to partition a set of size k into groups of size p if k is congruent to 1 modulo p . In our case, with $p = 2$ and k odd, this is equivalent to the statement that no perfect matching exists in an odd clique.

Likewise, via [Grigoriev, 2001, Definition 2] one checks that (3.6) constitutes a $PC_>$ proof; we refer the reader to Buss et al. [1999] for further discussion.

Applying Theorem 3.11 to (3.6), we find that $2k - 1 = \Omega(m) = \Omega(n)$, a contradiction when β is chosen small enough. Since $\tilde{S}(f) = \max f \leq (n-1)/2$ when f is associated with an odd set, we have $(1 - \varepsilon/(n-1))\tilde{C}(f) \geq \tilde{S}(f)$, which establishes an inapproximability ratio of $1 - \varepsilon/(n-1)$. \square

4 The Metric Traveling Salesperson Problem (TSP) revisited

In this section, we prove that a particular Lasserre SDP is optimal among all symmetric SDP relaxations for the asymmetric metric traveling salesperson problem on K_n . The *feasible solutions* of the problem are all permutations $\sigma \in S_n$. A permutation σ corresponds to the tour in K_n in which vertex i is the $\sigma(i)$ -th vertex visited. An *instance* \mathcal{I} of TSP is a set of non-negative distances $d_{\mathcal{I}}(i, j)$ for each edge $(i, j) \in K_n$, obeying the triangle inequality. The value of a tour σ is just the sum of the distances of edges traversed $\text{val}_{\mathcal{I}}(\sigma) = \sum_i d_{\mathcal{I}}(\sigma^{-1}(i), \sigma^{-1}(i+1))$. The *objective functions* are all the $\text{val}_{\mathcal{I}}$. Note that TSP is a *minimization* problem rather than a maximization problem, but the framework presented in Section 2 generalizes naturally to minimization problems by just flipping the inequalities. For approximation guarantees we will use $\tilde{S}(f) = \min f$ and $\tilde{C}(f) = \min f/\rho$ for some factor $\rho \geq 1$. Therefore instead of referring to a “ (\tilde{C}, \tilde{S}) -approximate formulation” we will refer to a “formulation within a factor ρ .”

The natural action of A_n on TSP is by permutation of vertices, which means here that A_n acts on S_n by composition from the left: $(\sigma_1 \cdot \sigma_2)(i) = \sigma_1(\sigma_2(i))$. Obviously, the problem TSP is A_n -symmetric.

The ring of real-valued functions on the set S_n of feasible solutions is isomorphic to $\mathbb{R}[\{x_{ij}\}_{\{i,j\} \in [n]^2}] / \langle Q_n \rangle_{\mathcal{I}}$, with x_{ij} being the indicator of $\sigma(i) = j$, and Q_n is the set of *TSP constraints*:

$$\begin{aligned} Q_n = & \left\{ \sum_{i \in [n]} x_{ij} - 1 \mid j \in [n] \right\} \cup \left\{ \sum_{j \in [n]} x_{ij} - 1 \mid i \in [n] \right\} \\ & \cup \{x_{ij}x_{ik} \mid i, j, k \in [n]\} \cup \{x_{ij}x_{kj} \mid i, j, k \in [n]\} \\ & \cup \{x_{ij}^2 - x_{ij} \mid i, j \in [n]\}. \end{aligned}$$

We emphasize that our description of the TSP constraints is different from the TSP polytope treated in Yannakakis [1991, 1988] and Fiorini et al. [2012, 2015]: the variables x_{ij} do not directly encode the edges of a Hamiltonian cycle but instead specify a permutation of n vertices, encoded as a perfect bipartite matching on $K_{n,n}$.

Following the framework presented in Lee et al. [2014], we define the Lasserre hierarchy for TSP as follows. The (dual of) the k -th level Lasserre SDP relaxation for a TSP instance \mathcal{I} is given by

$$\begin{aligned} & \text{Maximize } C \\ & \text{subject to } \text{val}_{\mathcal{I}} - C \simeq_{(Q_n, k)} \sum_p p^2 \quad \text{for some polynomials } p. \end{aligned}$$

We now state our main theorem regarding optimal SDP relaxations for TSP.

Theorem 4.1. *Suppose that there is some coordinate A_{2n} -symmetric SDP relaxation of size $r < \sqrt{\binom{n}{k}} - 1$ approximating TSP within some factor $\rho \geq 1$ for instances on $2n$ vertices. Then the $(2k-1)$ -level Lasserre relaxation approximates TSP within the factor of ρ on instances on n vertices.*

To prove [Theorem 4.1](#) there is an equivalent of [Proposition 3.3](#) we need for TSP tours, so that a small set of invariant functions depends only on the positions of a small number of indices. We start with the following proposition.

Proposition 4.2. *Let \mathcal{H} be an A_n -symmetric set of functions of size $\binom{n}{k}$ on the set of TSP tours $\sigma \in S_n$. Then for every $h \in \mathcal{H}$ there is a set $W \subseteq [n]$ of size less than k , such that $h(\sigma)$ depends only on the positions of the vertices in W in the tour σ , and the sign of σ as a permutation.*

Proof. For every $h \in \mathcal{H}$ we can apply [Lemma 3.2](#) to the stabilizer of h to obtain a subset $W \subseteq [n]$ of size at most k such that h is stabilized by $A([n] \setminus W)$. Thus for every tour σ , h is constant on the $A([n] \setminus W)$ -orbit of σ . This orbit is clearly determined by the positions of the vertices in W and, since $A([n] \setminus W)$ preserves signs, the sign of the permutation σ . \square

Next we give a reduction which allows us to eliminate the dependence of the functions $h \in \mathcal{H}$ on the sign of the permutation σ . In particular we encode every TSP tour σ on an n -vertex graph as some new tour $\Phi(\sigma)$ in a $2n$ -vertex graph, such that $\Phi(\sigma)$ is always an even permutation in S_{2n} .

Lemma 4.3. *Let \mathcal{I} be an instance of TSP on K_n . Then there exists an instance \mathcal{I}' of TSP on K_{2n} and an injective map $\Phi : S_n \rightarrow S_{2n}$ such that*

1. $\text{val}_{\mathcal{I}}(\sigma) = \text{val}_{\mathcal{I}'}(\Phi(\sigma))$ for all $\sigma \in S_n$.
2. For every tour $\tau \in S_{2n}$ there exists $\sigma \in S_n$ such that $\text{val}_{\mathcal{I}'}(\Phi(\sigma)) \leq \text{val}_{\mathcal{I}'}(\tau)$
3. For all $\sigma \in S_n$ the permutation $\Phi(\sigma)$ is even.

Proof. Given a TSP instance \mathcal{I} on K_n we construct a new instance \mathcal{I}' on K_{2n} as follows:

- For every vertex $i \in \mathcal{I}$ add a pair of vertices i and i' to \mathcal{I}' .
- For every distance $d(i, j)$ in \mathcal{I} add 4 edges all with the same distance $d(i, j) = d(i', j) = d(i, j') = d(i', j')$ to \mathcal{I}' .
- For every pair of vertices $i, i' \in \mathcal{I}'$ add an edge of distance zero, i.e. set $d(i, i') = 0$.

We will call a tour $\tau \in S_{2n}$ *canonical* if it visits i' immediately after i , i.e. $\sigma(i') = \sigma(i) + 1$. We will write T for the set of canonical tours in S_{2n} . It is easy to check using the triangle inequality that for every tour τ there is a canonical tour with no larger value. For every tour σ in \mathcal{I} define $\Phi(\sigma)$ to be the corresponding canonical tour in \mathcal{I}' . That is $\Phi(\sigma)(i) = 2\sigma(i) - 1$ and $\Phi(\sigma)(i') = 2\sigma(i)$. Note that $\Phi : S_n \rightarrow S_{2n}$ is an injective map whose image is all of T . By construction we have:

$$\text{val}_{\mathcal{I}}(\sigma) \equiv \text{val}_{\mathcal{I}'}(\Phi(\sigma))$$

which proves property (1). Property (2) follows from the fact that every tour $\tau \in S_{2n}$ has a canonical tour with no larger value, and that T is the image of Φ .

For property (3), note that every canonical tour is an even permutation. To see why, suppose $\sigma \in S_n$ is given by $\sigma = (i_1, j_1)(i_2, j_2), \dots, (i_m, j_m)$ where (i, j) denotes the permutation that swaps i and j . Then $\Phi(\sigma) = (i_1, j_1)(i'_1, j'_1), \dots, (i_m, j_m)(i'_m, j'_m)$ is comprised of $2m$ swap permutations, and is therefore even. \square

The last ingredient we need is a version of [Theorem 3.9](#) for the TSP.

Theorem 4.4. *If F is a multilinear polynomial whose monomials are partial matchings on $K_{n,n}$ and $F \in \langle Q_n \rangle_I$, then $F \simeq_{(Q_n, 2 \deg F - 1)} 0$.*

Because Q_n is so similar to \mathcal{P}_n , it should come as no surprise that the proof of the above theorem is extremely similar to the proof of [Theorem 3.9](#). We include the full proof for completeness, but defer it to [Section 4.1](#). We now have all the tools necessary to prove [Theorem 4.1](#).

Proof of Theorem 4.1. First let I be an instance of TSP on K_n . Use [Lemma 4.3](#) to construct a TSP instance I' on K_{2n} and the corresponding map Φ . Now assume we have an arbitrary A_{2n} -symmetric SDP relaxation of size $d < \sqrt{\binom{2n}{k}} - 1$ for TSP on K_{2n} . By [Lemma 2.3](#) there is a corresponding A_{2n} -symmetric family of functions \mathcal{H}' of size $\binom{d+1}{2}$ such that whenever $\min_\tau \text{val}_{I'}(\tau) \geq \tilde{S}(\text{val}_{I'})$ we have:

$$\text{val}_{I'}(\tau) - \tilde{C}(\text{val}_{I'}) \equiv \sum_j h_j(\tau)^2 + \mu_{I'} \quad \text{where } h_j \in \langle \mathcal{H}' \rangle \text{ and } \mu_{I'} \geq 0.$$

Let $h' \in \mathcal{H}'$. By [Proposition 4.2](#) $h'(\tau)$ depends only on some subset W' of size at most k , and possibly on the sign of τ .

Now we restrict the above relaxation to the image of Φ . By [Lemma 4.3](#) this does not change the optimum. Using the fact that $\text{val}_I(\sigma) \equiv \text{val}_{I'}(\Phi(\sigma))$ and setting $\mu_I = \mu_{I'}$ then gives rise to a new relaxation where whenever $\min_\sigma \text{val}_I(\sigma) \geq \tilde{S}(\text{val}_I)$ we have:

$$\text{val}_I(\sigma) - \tilde{C}(\text{val}_I) \equiv \sum_j h_j(\Phi(\sigma))^2 + \mu_I \quad \text{where } h_j \in \langle \mathcal{H}' \rangle \text{ and } \mu_I \geq 0$$

as $\tilde{S}(\text{val}_I) = \tilde{S}(\text{val}_{I'})$ and $\tilde{C}(\text{val}_I) = \tilde{C}(\text{val}_{I'})$ by [Lemma 4.3](#). Next for each $h' \in \mathcal{H}'$ define $h : S_n \rightarrow \mathbb{R}$ by $h(\sigma) = h'(\Phi(\sigma))$. Since $\Phi(\sigma)$ is even, we then have that each h depends only on the position of some subset $W \subseteq [n]$ of size at most k . Such a function can be written as a degree- k polynomial p in the variables x_{ij} so that $p(x^\sigma) \equiv f(\sigma)$ on the vertices of $P_{TSP}(n)$. Now by [Theorem 4.4](#) we have that $p \simeq_{(Q_n, 2k-1)} h$. Since $\mu_I \geq 0$ it is clearly the square of a (constant) polynomial, and we conclude that whenever $\min_\sigma \text{val}_I(\sigma) \leq \tilde{S}(\text{val}_I)$ we have:

$$f_I(x) - \min f_I / \rho \simeq_{(Q_n, 2k-1)} \sum_p p(x)^2$$

which is precisely the statement that the $(2k - 1)$ -level Lasserre relaxation for $P_{TSP}(n)$ is a ρ -approximation. \square

4.1 Low-degree certificates for tour ideal membership

In this section we prove [Theorem 4.4](#) showing that every degree- d polynomial identically zero over *TSP tours* is congruent to 0 within degree $O(d)$.

Note that any partial tour τ can be thought of as a partial matching M in $K_{n,n}$, namely if $\tau(i) = j$, then M includes the edge (i, j) . Because of this, it will come as no surprise that the proof proceeds in a very similar manner to [Section 3.4](#), and hereafter we shall always refer to partial matchings on $K_{n,n}$ rather than on K_n .

For a partial matching M , let $x_M := \prod_{e \in M} x_e$ denote the product of edge variables for the edges in M . The first step is to reduce every polynomial to a linear combination of the x_M .

Lemma 4.5. *For every polynomial F there is a polynomial F' with $\deg F' \leq \deg F$ and $F \simeq_{(Q_n, \deg F)} F'$, where all monomials of F have the form x_M for some partial matching M .*

Proof. It is enough to prove the lemma when F is a monomial: $F = \prod_{e \in A} x_e^{k_e}$ for a set $A \subseteq E[K_{n,n}]$ of edges with multiplicities $k_e \geq 1$. From $x_e^2 \simeq_2 x_e$ it follows that $x_e^k \simeq_k x_e$ for all $k \geq 1$, hence $F \simeq_{\deg F} \prod_{e \in A} x_e$, proving the claim if A is a partial matching. If A is not a partial matching, then there are distinct $e, f \in A$ with a common vertex, hence $x_e x_f \simeq_2 0$ and $F \simeq_{\deg F} 0$. \square

The rest of the proof proceeds identically to [Theorem 3.9](#), but we let the symmetric group act on polynomials slightly differently. If $K_{n,n} = U_n \cup V_n$ is the bipartite decomposition of $K_{n,n}$, then we only let the permutation group act on the labels of vertices of U_n , i.e. $\sigma x_{(a,b)} = x_{(\sigma(a),b)}$. We show that under this action, symmetrized polynomials are congruent to a constant, which can again be seen in the same sequence of lemmas:

Lemma 4.6. *For any partial matching M on $2d$ vertices and a vertex $a \in U_n$ not covered by M , we have*

$$x_M \simeq_{(Q_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, u\} \\ u \in V_n \setminus (M \cap V_n)}} x_{M_1}. \quad (4.1)$$

Proof. We use the generators $\sum_v x_{av} - 1$ to add variables corresponding to edges at a , and then use $x_{av} x_{bv}$ to remove monomials not corresponding to a partial matching:

$$x_M \simeq_{(Q_n, d+1)} x_M \sum_{v \in V_n} x_{av} \simeq_{(Q_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, v\} \\ v \in V_n \setminus (M \cap V_n)}} x_{M_1}.$$

\square

This leads to a similar congruence using all containing matchings of a larger size:

Lemma 4.7. *For any partial matching M of $2d$ vertices and $d \leq k \leq n$, we have*

$$x_M \simeq_{(Q_n, k)} \frac{1}{\binom{n-d}{k-d}} \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'} \quad (4.2)$$

Proof. We use induction on $k - d$. The start of the induction is when $k = d$, when the sides of Equation (4.2) are equal.

If $k > d$, let $a \in U_n$ be a fixed vertex not covered by M . Applying [Lemma 4.6](#) to M and a followed by the inductive hypothesis gives:

$$x_M \simeq_{(Q_n, d+1)} \sum_{\substack{M_1 = M \cup \{a, u\} \\ u \in V_n \setminus (M \cap V_n)}} x_{M_1} \simeq_{(Q_n, k)} \frac{1}{\binom{n-d-1}{k-d-1}} \sum_{\substack{M' \supset M_1 \\ |M'|=k \\ M_1 = M \cup \{a, u\} \\ u \in V_n \setminus (M \cap V_n)}} x_{M'}.$$

Averaging over all vertices $a \in U_n$ not covered by M , we obtain

$$x_M \simeq_{(Q_n, k)} \frac{1}{n-d} \frac{1}{\binom{n-d-1}{k-d-1}} \sum_{\substack{M' \supset M_1 \\ |M'|=k \\ M_1 = M \cup \{a, u\} \\ a \in U_n \setminus (M \cap U_n) \\ u \in V_n \setminus (M \cap V_n)}} x_{M'} = \frac{1}{n-d} \frac{1}{\binom{n-d-1}{k-d-1}} (k-d) \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'} = \frac{1}{\binom{n-d}{k-d}} \sum_{\substack{M' \supset M \\ |M'|=k}} x_{M'}.$$

\square

Corollary 4.8. *For any polynomial F , there is a constant c_F with $\sum_{\sigma \in S_n} \sigma F \simeq_{(Q_n, \deg F)} c_F$.*

Proof. In view of Lemma 4.5, it is enough to prove the claim for $F = x_M$ for some partial matching M on $2k$ vertices, which is an easy application of Lemma 4.7 with $d = 0$:

$$\sum_{\sigma \in S_n} \sigma x_M = (n - k)! \sum_{M' : |M'|=k} x_{M'} \simeq_k (n - k)! \binom{n}{k}.$$

□

The next lemma will allow us to apply induction:

Lemma 4.9. *If L is a polynomial with $L \simeq_{(Q_{n-2}, d)} 0$ and a, b are the additional vertices in Q_n then $Lx_{ab}x_{ba} \simeq_{(Q_n, d+2)} 0$.*

Proof. It is enough to prove the claim when L is from Q_{n-2} . For $L = x_e^2 - x_e$, $L = x_{uv}x_{uw}$, and $L = x_{uv}x_{wv}$ the claim is trivial, as then $L \in Q_n$. The remaining cases are (1) $L = \sum_{u \in U_{n-2}} x_{uv} - 1$ for some $v \in V_{n-2}$, and (2) $L = \sum_{v \in V_{n-2}} x_{uv} - 1$ for some $u \in U_{n-2}$. We only deal with the first case, as the second one is analogous. Then

$$Lx_{ab}x_{ba} = \left(\sum_{u \in U_n} x_{uv} - 1 \right) x_{ab}x_{ba} - x_{av}x_{ab}x_{ba} - x_{bv}x_{ab}x_{ba} \simeq_{(Q_n, d+1)} 0.$$

□

We are now ready to prove Theorem 4.4.

Proof of Theorem 4.4. We use induction on the degree d of F . The case $d = 0$ is obvious, as then clearly $F = 0$. (Note that \simeq_{-1} is just equality.) The case $d = 1$ rephrased means that the affine space spanned by the characteristic vectors of all perfect matchings is defined by the $\sum_v x_{uv} - 1$ for all vertices u . This follows again from Edmonds's description of the perfect matching polytope by linear inequalities in Edmonds [1965] (valid for any graph in addition to K_{2n} and $K_{n,n}$).

For the case $d \geq 2$ we first prove the following claim:

Claim. If $F \in \langle Q_n \rangle_I$ is a degree- d polynomial and $\sigma \in S_n$ is a permutation of vertices, then

$$F \simeq_{(Q_n, 2d-1)} \sigma F.$$

We use induction on the degree. If $d = 0$ or $d = 1$ the claim follows from the corresponding cases $d = 0$ and $d = 1$ of the theorem. For $d \geq 2$ it is enough to prove the claim when σ is a transposition of two vertices a and u . Note that in $F - \sigma F$ all monomials which do not contain an x_e with e incident to a or u on the left cancel:

$$F - \sigma F = \sum_{e: e=(a,r) \text{ or } e=(u,r)} L_e x_e \tag{4.3}$$

where each L_e has degree at most $d - 1$. We now show that every summand is congruent to a sum of monomials containing edges incident to both a and u on the left. For example, for $e = \{a, b\}$ in (4.3), we apply the generator $\sum_v x_{uv} - 1$ to find:

$$L_{ab}x_{ab} \simeq_{d+1} L_{ab}x_{ab} \sum_v x_{uv} \simeq_{d+1} \sum_v L_{ab}x_{ab}x_{uv}.$$

Therefore

$$F - \sigma F \simeq_{d+1} \sum_{bv} L'_{bv} x_{ab} x_{uv}$$

for some polynomials L'_{bv} of degree at most $d-1$. We may assume that L'_{bv} does not contain variables x_e with e incident to a, u on the left or b, v on the right, as these can be removed using generators like $x_{ab}x_{ac}$ or $x_{ab}^2 - x_{ab}$. Moreover, since F is zero on all perfect matchings, it can be checked that L'_{bv} is zero on all perfect matchings containing $\{a, b\}$ and $\{u, v\}$. By induction, $L'_{bv} \simeq_{(Q_{n-4}, 2d-3)} 0$ (identifying K_{n-4} with the graph $K_n \setminus \{a, b, u, v\}$), from which $L'_{bv} \simeq_{(Q_n, 2d-1)} 0$ follows by two applications of [Lemma 4.9](#). (The special case $a = v, b = u$ is also handled by induction and one application of [Lemma 4.9](#).) This concludes the proof of the claim.

We now apply the claim followed by [Corollary 4.8](#):

$$F \simeq_{2d-1} \frac{1}{n!} \sum_{\sigma \in S_n} \sigma F \simeq_d \frac{c_F}{n!}$$

for a constant c_F . As $F \in \langle Q_n \rangle_I$, it must be that $c_F = 0$, and therefore $F \simeq_{2d-1} 0$. \square

References

- Abbas Bazzi, Samuel Fiorini, Sebastian Pokutta, and Ola Svensson. No small linear program approximates vertex cover within a factor $2 - \varepsilon$. In *Proc. FOCS*, pages 1123–1142, 2015. doi: [10.1109/FOCS.2015.73](https://doi.org/10.1109/FOCS.2015.73). 2
- Gábor Braun and Sebastian Pokutta. An algebraic approach to symmetric extended formulations. In *Proc. ISCO*, pages 141–152. Springer Berlin Heidelberg, 2012. ISBN 978-3-642-32147-4. doi: [10.1007/978-3-642-32147-4_14](https://doi.org/10.1007/978-3-642-32147-4_14). 6
- Gábor Braun and Sebastian Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. In *Proc. SODA*, pages 837–846, 2015a. doi: [10.1137/1.9781611973730.57](https://doi.org/10.1137/1.9781611973730.57). 5
- Gábor Braun and Sebastian Pokutta. The matching problem has no fully polynomial size linear programming relaxation schemes. *IEEE Transactions on Information Theory*, 61(10):5754–5764, 2015b. doi: [10.1109/TIT.2015.2465864](https://doi.org/10.1109/TIT.2015.2465864). 5
- Gábor Braun, Samuel Fiorini, Sebastian Pokutta, and David Steurer. Approximation limits of linear programs (beyond hierarchies). In *Proc. FOCS*, pages 480–489, 2012. 2
- Gábor Braun, Samuel Fiorini, Sebastian Pokutta, and David Steurer. Approximation limits of linear programs (beyond hierarchies). *Math. Oper. Res.*, 40(3):756–772, 2015a. doi: [10.1287/moor.2014.0694](https://doi.org/10.1287/moor.2014.0694). 2
- Gábor Braun, Sebastian Pokutta, and Daniel Zink. Inapproximability of combinatorial problems via small LPs and SDPs. In *Proc. STOC*, pages 107–116, 2015b. doi: [10.1145/2746539.2746550](https://doi.org/10.1145/2746539.2746550). 3, 5
- Mark Braverman and Ankur Moitra. An information complexity approach to extended formulations. In *Proc. STOC*, pages 161–170. ACM, 2013. ISBN 978-1-4503-2029-0. doi: [10.1145/2488608.2488629](https://doi.org/10.1145/2488608.2488629). 2
- Jop Briët, Daniel Dadush, and Sebastian Pokutta. On the existence of 0/1 polytopes with high semidefinite extension complexity. In *Proc. ESA*, pages 217–228. Springer Berlin Heidelberg, 2013. 1
- Jop Briët, Daniel Dadush, and Sebastian Pokutta. On the existence of 0/1 polytopes with high semidefinite extension complexity. *Math. Program.*, 153(1):179–199, 2015. doi: [10.1007/s10107-014-0785-x](https://doi.org/10.1007/s10107-014-0785-x). 1
- Sam Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. In *Proc. STOC*, pages 547–556, 1999. 11
- Siu On Chan, James R. Lee, Prasad Raghavendra, and David Steurer. Approximate constraint satisfaction requires large LP relaxations. In *Proc. FOCS*, pages 350–359, 2013. 2
- John D. Dixon and Brian Mortimer. *Permutation groups*. Springer Verlag, 1996. 6
- Jack Edmonds. Maximum matching and a polyhedron with 0,1-vertices. *J. Res. Nat. Bur. Standards Sect. B*, 69B:125–130, 1965. ISSN 0160-1741. 8, 15

Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds. In *Proc. STOC*, pages 95–106, 2012. [2](#), [11](#)

Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Exponential lower bounds for polytopes in combinatorial optimization. *J. Assoc. Comput. Mach.*, 62(2):17, 2015. doi: [10.1145/2716307](#). [2](#), [11](#)

Michel X. Goemans. Smallest compact formulation for the permutohedron. *Math. Program.*, 153(1): 5–11, 2015. doi: [10.1007/s10107-014-0757-1](#). [2](#)

Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42: 1115–1145, 1995. doi: [10.1145/227683.227684](#). [1](#)

João Gouveia, Pablo A Parrilo, and Rekha R Thomas. Lifts of convex sets and cone factorizations. *Math. Oper. Res.*, 38(2):248–264, May 2011. [1](#)

Dima Grigoriev. Linear lower bound on degrees of positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1):613–622, 2001. [2](#), [10](#), [11](#)

Volker Kaibel, Kanstantsin Pashkovich, and Dirk Oliver Theis. Symmetry matters for the sizes of extended formulations. In *Proc. IPCO*, pages 135–148, 2010. doi: [10.1007/978-3-642-13036-6_11](#). [1](#), [6](#)

James R. Lee, Prasad Raghavendra, David Steurer, and Ning Tan. On the power of symmetric LP and SDP relaxations. In *Proc. CCC*, pages 13–21, 2014. [2](#), [3](#), [6](#), [11](#)

James R. Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *Proc. STOC*, pages 567–576, 2015. doi: [10.1145/2746539.2746599](#). [1](#), [2](#)

Kanstantsin Pashkovich. Tight lower bounds on the sizes of symmetric extensions of permutohedra and similar results. *Math. Oper. Res.*, 39(4):1330–1339, 2014. doi: [10.1287/moor.2014.0659](#). [2](#)

Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proc. STOC*, pages 263–272, 2014. [1](#), [2](#)

Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discrete Math.*, 3: 411–430, 1990. doi: [10.1137/0403036](#). [2](#)

Lieven Vandenberghe and Stephen Boyd. Semidefinite programming. *SIAM Rev.*, 38:49–95, 1996. [1](#)

Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs (extended abstract). In *Proc. STOC*, pages 223–228, 1988. [1](#), [2](#), [6](#), [11](#)

Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *J. Comput. System Sci.*, 43(3):441–466, 1991. ISSN 0022-0000. doi: [10.1016/0022-0000\(91\)90024-Y](#). [1](#), [2](#), [6](#), [11](#)